

AD-A221 468

SRI International



DTIC FILE COPY

(1)

ANALYSIS OF ROUTING STRATEGIES FOR PACKET RADIO NETWORKS

SRNTN 14

August 1984

By: J. J. Garcia-Luna-Aceves
Nachum Shacham

Telecommunications Sciences Center
Computer Science and Technology Division

Prepared for:

Defense Advanced Research Projects Agency
1400 Wilson Boulevard
Arlington, Virginia 22209

Attention: Dr. Barry M. Leiner

SRI Project 5732

Effective Date: March 17, 1983
Expiration Date: May 15, 1984

Sponsored by
The Defense Advanced Research Projects Agency (DoD)
ARPA Order No. 4570
Under Contract No. MDA-903-83-C-0155 issued by
Department of Army, Defense Supply Service-Washington,
Washington, D.C. 20310

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency, the Army, or the United States Government.

This work was also sponsored by the U.S. Army Research Office under Contract DAAG29-84-K-0042

DTIC
ELECTE
MAY 15 1990
Dcs D

DO NOT REMOVE
ZUAAAAAA2594942

DISTRIBUTION STATEMENT A

Approved for public release
Distribution Unlimited

333 Ravenswood Ave. • Menlo Park, CA 94025
415 326-6200 • TWX 910-373-2046 • Telex 334-486

90 05 14 090

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Analysis of Routing Strategies for Packet Radio Networks		5. TYPE OF REPORT & PERIOD COVERED Survivable Radio Networks Temporary Note
		6. PERFORMING ORG. REPORT NUMBER 5732-SRNTN #14
7. AUTHOR(s) Jose J. Garcia Luna Aceves Nachum Shacham		8. CONTRACT OR GRANT NUMBER(s) MDA903-83-C-0155
9. PERFORMING ORGANIZATION NAME AND ADDRESS SRI International 333 Ravenswood Avenue Menlo Park, CA 94025		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS Program Code Number P62708E
11. CONTROLLING OFFICE NAME AND ADDRESS Defense Advanced Research Projects Agency 1400 Wilson Boulevard Arlington, Virginia 22209		12. REPORT DATE
		13. NUMBER OF PAGES
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS (of this report) UNCLASSIFIED
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Distribution Unlimited		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Large packet radio networks, Clusters, Global routing node (GRN), Routing table, Routing algorithms		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) In this paper we present a comparative analysis of network routing schemes for large-scale packet radio networks (PRNET) with mobile nodes. This analysis is aimed at determining two things: (1) the time required to obtain consistent routing tables in all the nodes of a PRNET after topological changes, and (2) the optimality of the routing decisions made by a node in terms of the length of the paths chosen to remote destina- tions. We make minimal assumptions about the routing protocols used and		

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

the topology of the PRNET to obtain upper bounds on the length of shortest paths and on the time required for routing-table updating after topological changes in PRNETs with hierarchical organizations. Such bounds are used to analyze the optimization of network topology and as guidelines for the design of hierarchical routing schemes for large PRNETs. Such results extend previous ones obtained by Kamoun [KAMO-76], Hagouel [HAGO-83], and Baratz and Jaffe [BARA-83] for land-based networks.

Accession For	
NTIS CR&I	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution /	
Availability Codes	
Dist	Avail and/or Special
A-1	



UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

1. INTRODUCTION

The dynamic determination of optimum routes between nodes is fundamental in the operation of multihop packet-radio networks (PRNET), but may become very costly for PRNETs with several mobile nodes. A number of routing schemes have been proposed in the past to cope with this problem, and can be classified as centralized, distributed, and hierarchical. In the centralized scheme, also called station-mode routing [KAHN-78], a single node (called the station) ascertains the best path between each pair of nodes and upon request sends the requisite routing information to nodes in the PRNET. This routing strategy would be unacceptable for large PRNETs because of its inherent vulnerabilities [GAFN-81; WEST-82]. In a PRNET using the fully distributed scheme (also called stationless mode [WEST-82]), which we shall call a *flat* PRNET, all the nodes participate as peers in the same distributed algorithm to determine dynamically the best path to every node. Finally, a number of hierarchical routing schemes have been proposed for the management of routing information in large PRNETs [BAKE-81, 83; MACG-82; NILS-80; SHAC-84a]. The main idea of such schemes is to allow each node to maintain exact routing information regarding nodes very close to it, and less detailed information regarding nodes farther away from it. The objective of doing so is to obtain a reasonable compromise among the size of routing tables, number of updates required to maintain such tables, and the speed with which updates are propagated.

In this paper, we analyze the performance of the hierarchical routing strategy for large, mobile PRNETs previously proposed by Shacham and Klemba [SHAC-84a, 84c], and compare it with other hierarchical and fully distributed routing strategies. We focus on two main performance figures in our analysis: (1) the time required to obtain consistent routing tables at all the nodes of a PRNET after topological changes, and (2) the length of the paths that can be obtained with different routing strategies. We chose to analyze worst-case network performance, rather than average network performance, to be able to obtain as general PRNET design guidelines as possible by making minimum assumptions about the topological characteristics of the PRNET.

In Section 2 we describe the hierarchical routing scheme proposed by Shacham and Klemba, and in Section 3 we compare its performance with those of other schemes. In doing so, we extend the results presented by Hagouel [HAGO-83] and by Baratz and Jaffe [BARA-83] on path lengths in hierarchical networks; our results on path lengths can also be applied to networks with point-to-point links. In Section 4 we address the optimization of network organization with respect to the quality of routing in PRNETs using the upper bounds on path lengths obtained in section 3. Finally, in Section 5 we discuss the results presented in sections 3 and 4, and outline how our results can be applied to the design of PRNETs.


2. A HIERARCHICAL ROUTING STRATEGY FOR LARGE PRNETS

2.1 BASIC NETWORK ORGANIZATION


The routing scheme proposed by Shacham and Klemba [SHAC-84a, 84c] is similar to the multistation scheme [BAKE-81], and is based on three major premises:

- (1) The nodes of the PRNET are organized into m levels of clusters (where $m \geq 1$) to reduce the length of routing tables. Nodes represent clusters at Level 0, a group of nodes is a cluster at Level 1 (called 1-cluster), and a cluster at Level k (called k -cluster) is the union of clusters at Level $k - 1$.
- (2) The updating of routing information among clusters is carried out on an *event-driven* basis (i.e., immediately after every topological change), but by only a few nodes called *global-routing nodes* or *GRN* (one per 1-cluster). This is done in the hope of accelerating the dissemination of updates that affect a large number of nodes, while update traffic levels are kept down.
- (3) Each node participates in two parallel updating procedures. One procedure updates routing information about other nodes that are close by; the other updates routing information about distant nodes organized into clusters and is controlled by *GRNs*.

Figure 1 illustrates a PRNET organized into three levels of clusters; the links between nodes of the figure indicate radio connectivity. Every node must be affiliated with at least one 1-cluster to communicate with the other nodes in the PRNET. Those nodes that have radio connectivity with nodes in different 1-clusters are called *boundary nodes*. Two 1-clusters that have *at least* one *boundary node* in common (i.e., one that is affiliated with both clusters) are said to *overlap*. In contrast, two 1-clusters connected by *boundary nodes* that are only affiliated to any one of them are said to be adjacent disjoint 1-clusters. Note that clusters at levels 2 and above never overlap (to be explained later). The difference between overlapping and adjacent disjoint 1-clusters can be appreciated by observing Figure 1(a). Cluster A.1 and Cluster A.2 overlap because Node a is affiliated with both clusters. In contrast, Cluster A.1 and A.3 are disjoint but adjacent because nodes a and b have radio connectivity with each other but belong to only one of the two 1-clusters. The procedure by which a node becomes affiliated with a 1-cluster is not addressed in this paper.



(a)



(b)

AT NODE 1:

NRT			GRT			
DEST NODE	NEXT NODE	DISTANCE	DEST CLUSTER	NEXT GRN	NEXT BN	DISTANCE
A.1.a	A.1.a	1	A.1	A.2.9	A.1.a	0
A.1.c	A.1.a	2	A.2	A.3.2	A.1.a	1
A.1.1		0	A.3			1
			A			0
			B	A.2.9	A.1.a	2
			C	A.2.9	A.1.a	5

(c)

FIGURE 1 A G-NETWORK

Each node maintains two routing tables:

- (1) A *node-level routing table (NRT)*, which contains routing information about nodes in the same 1-cluster with which the node is affiliated.
- (2) A *group-level routing table (GRT)*, which contains routing information about the clusters in the same higher-level clusters to which the node belongs.

The nodes within a given 1-cluster update their *NRTs* as if they constituted a small, flat PRNET. The *GRNs* of the PRNET are organized as a *virtual network*; the nodes of this network are the *GRNs*, and a link is defined between two *GRNs* if and *only* if there is radio connectivity among boundary nodes of their respective 1-clusters. The messages exchanged among *GRNs* are forwarded through multihop paths formed by node-to-node links; simple nodes in those paths simply forward such messages towards the destination *GRNs* and retransmit the messages as necessary to ensure reliable transmission. Hence, the network of *GRNs* constitutes a *virtual point-to-point network*; Figure 1(b) illustrates the virtual network of *GRNs* for the PRNET of Figure 1(a). The virtual network of *GRNs* utilizes the routing scheme proposed by Kamoun and Kleinrock [KAMO-76] to update the *GRTs* of *GRNs*. A more detailed description of the contents of *NRTs* and *GRTs*, together with the procedures followed to update them, is provided in the next two subsections.

2.2 NRT UPDATE

In *small*, single-channel PRNETs with *flat* organizations, a node can receive only one collision-free message at a time from the radio channel, nodes can be highly mobile, and link quality may change fairly often. In such PRNETs, it appears that periodic routing-table update algorithms based on next-node tables* and with no retransmission of updates, such as the tier-routing algorithm [WEST-82], are a good choice for such networks, and is, therefore, the type of algorithm used in our scheme. Simple nodes, boundary nodes, and *GRNs* within the same 1-cluster participate as peers in a periodic routing update algorithm on the basis of their *NRTs*. The *NRT* of a given node contains an entry for each node in its 1-cluster; each entry specifies the next node to a given 1-cluster destination and the length (i.e., number of links between nodes) of the shortest path to that destination. *NRTs* are updated by means of *node-level updates (NLUs)* transmitted periodically and without retransmissions within a 1-cluster. An *NLU* contains all the entries in the *NRT* and *GRT* of a node; the rationale for including the content *GRTs* in the *NLUs* is explained below.

* A next-node table specifies summary routing information consisting of the next node and the length of the minimum path to every node in the PRNET.

Note that boundary nodes and *GRNs* must maintain an entry in their *NRTs* for each node in all the 1-clusters with which they are affiliated. Whenever either a simple node or a *GRN* receives *NLU*s from boundary nodes referring to 1-clusters with which it is not affiliated, it simply ignores them. This guarantees that information in *NRTs* is not propagated across clusters' boundaries.

2.3 GRT UPDATE

Because of the point-to-point nature of a virtual network of *GRNs*, it is possible to employ a reliable, event-driven algorithm to maintain consistent *GRTs*, i.e., an algorithm in which updates are sent whenever topological changes occur and in which reliable transmission of updates is ensured. More specifically, the network of *GRNs* is organized in m cluster levels by means of the Kamoun-Kleinrock scheme, which implies that

- (1) Clusters of *GRNs* (clusters at levels 2 and above) are disjoint.
- (2) All *GRNs* in the network of *GRNs* participate as peers in the same algorithm to update the entries of their *GRTs*.
- (3) The *GRT* of a *GRN* contains $m - 1$ *j*-subtables ($1 \leq j \leq m - 1$). A *j*-subtable contains entries for all *j*-clusters within the *GRN's* ($j + 1$)-cluster. Each such entry specifies: (a) the destination *j*-cluster; (b) the next *GRN* in the chain of *GRNs* to that destination; (c) at least one boundary node towards the next *GRN*; (d) the number of *GRN*-to-*GRN* hops in that chain.

Entry (b) above permits routing of messages from a node to remote clusters through adjacent clusters, while entry (c) allows a node to route messages to boundary nodes in its own 1-cluster towards remote clusters. The distance from a *GRN* to its own *k*-cluster is set to 0, while the length of the shortest path from a *GRN* to a remote *k*-cluster equals one plus the minimum of the shortest path lengths reported by its neighbor *GRNs* for that destination. Figure 1 (c) illustrates the content of the *GRT* and *NRT* for Node 1 of Figure 1 (a).

Cluster-level updates (CLUs) from a *GRN* are transmitted reliably on an event-driven basis to all its neighbor *GRNs* in the same way in which updates are transmitted among nodes in the scheme proposed by Kamoun and Kleinrock [KAMO-76]. Each *CLU* contains entries consisting of the identifier of a destination *j*-cluster, the next *GRN* in the chain towards it, and the number of *GRN*-to-*GRN* hops in that chain. Each such entry corresponds to an entry of a *GRN's GRT* that was updated because of a change in connectivity with neighbor *GRNs*, or because of *CLUs* received from other *GRNs*. A *GRN* sends a *CLU* to a neighbor *GRN* containing only that routing information that refers to common destinations in their *GRTs*. Hence, if two *GRNs*, x and y , are in the same *k*-cluster, but in different $(k - 1)$ -clusters, *GRN* x sends *CLUs* to *GRN* y that refer to

clusters at levels equal to or larger than k .

As we have said, each node stores a *GRT* with the information necessary to route packets across 1-clusters. However, the *GRNs* are the only ones that can initiate the update of *GRTs* by exchanging *CLUs*; simple nodes or boundary nodes that receive *CLUs* simply forward them without any further processing. When *GRNs* generate *CLUs*, they distribute the completely updated *GRTs* to simple nodes and boundary nodes as part of the next *NLU*s transmitted periodically by the *GRNs* of 1-clusters. Simple nodes and boundary nodes are capable of updating their *GRTs* because each *NLU* contains the entire *GRT* and *NRT* of the transmitting node; hence, a *GRN* that updates its *GRT* communicates such updates to the rest of the nodes in its 1-cluster in the next *NLU* (transmitted periodically). Note that a complete *GRT* and *NRT* must be included in each *NLU* because *NLU*s may be lost and no *NLU* is retransmitted.

2.4 ROUTING AMONG NODES IN DIFFERENT CLUSTERS

Consider two simple nodes who lie within the same $(k + 1)$ -cluster, but in different k -clusters: a simple Node a in k -cluster A and Node b in k -cluster B . The only information that Node a has to route messages to Node b consists of the next *GRN* towards k -cluster B and one or more boundary nodes within Node a 's 1-cluster towards that *GRN*. Accordingly, nodes carry out routing as follows:

- (1) Node a looks up its *GRT* to obtain the next *GRN* towards the destination cluster; the entry in its *GRT* provides *at least* one boundary node towards that destination.
- (2) Node a looks up its *NRT* for a next node towards that boundary node.
- (3) The nodes in the path from a to b perform the same type of procedure.

Boundary nodes must move packets across boundary nodes. As we have stated previously, boundary nodes of overlapping clusters maintain routing information about all the nodes in the 1-clusters to which they belong; hence, they can forward messages across 1-cluster boundaries as any simple node. One way to support routing across nonoverlapping 1-clusters is for boundary nodes to add an entry in their *NRT* for each adjacent boundary node; such entries would specify that the adjacent boundary nodes are 1 hop away (i.e., they are the next nodes towards themselves).

3. PERFORMANCE OF THE PROPOSED SCHEME

In this section we compare the hierarchical routing scheme being proposed with other schemes. Specifically, we consider the following cases:


- (1) A PRNET organized according to the scheme described in the previous section, which we shall refer to as a *G-network*.
- (2) A flat PRNET
- (3) A PRNET in which *all* the nodes are organized according to the Kamoun-Kleinrock scheme. We shall refer to this type of PRNET as a *K-network*

Figure 2 illustrates the structure of K-networks using the same PRNET depicted in Figure 1 as a G-network. An examination of both figures shows that the basic difference among K-networks and G-networks is the metric assumed to measure the distance from a node to any other node in a remote cluster. In the G-network of Figure 1, the distance between Node 1 and any destination in a remote j -cluster ($j \geq 1$) is the number of GRN-to-GRN hops from Node 1's 1-cluster to any 1-cluster in the remote j -cluster; for instance, the distance from Node 1 to any destination within 2-cluster C is *five hops*. In contrast, in the K-network of Figure 2, the distance between a node and a destination in a remote cluster (i.e., the highest-level cluster that the source and the destination do not share) is measured by the shortest path (in node-to-node hops) to a boundary node in such a remote cluster. For instance, the distance from Node 1 to any node in 2-cluster C is *twelve node-to-node hops*. Note that if the PRNET had a flat organization, Node 1 would have to know the shortest distance to all the other 26 nodes in the PRNET.

3.1 MAXIMUM PATH LENGTHS—QUALITY OF ROUTING

Organizing a PRNET into hierarchies reduces the amount of information possessed by each node about the topology of the network. Accordingly, the routing decisions made by a hierarchical-network node may not yield the best possible routes; in this subsection we quantify this effect. We obtain the ratio of the worst-case path lengths that can be obtained in G-networks and K-networks with respect to the optimum path lengths obtained in flat networks. We must point out that the worst cases for G-paths and K-paths described here can in fact be attained, as will be apparent from the derivations that follow.

We shall assume throughout this subsection that the routing tables (*NRTs* and *GRTs*) of network nodes are correct and that clusters do not overlap. Furthermore, we consider that links are bidirectional and that internodal distances are measured in number of hops. We shall refer to the shortest path between two nodes of a flat PRNET as a *flat path*.



(a)

AT NODE 1:

DESTINATION	NEXT NODE	DISTANCE
a		1
c	A.1.a	2
1	A.1.a	0
A.1		0
A.2	A.1.a	2
A.3	A.1.a	2
A		0
B	A.1.a	4
C	A.1.a	12

(b)

FIGURE 2 A K-NETWORK

3.1.1 K-Networks

Hagouel has shown that, in the worst case, the shortest paths obtained between two nodes of different m -clusters in a k -level network ($m \leq k$) based on routing tables constructed according to the Kamoun-Kleinrock scheme can be $2^m - 1$ times longer than the minimum paths that would be obtained in a flat network [HAGO-83]. Baratz and Jaffe [BARA-83] have shown the same result for the case in which $m = 2$. This result was obtained under the assumption that the actual minimum path between any two nodes in the network must always lie within a 1-cluster common to both nodes. However, as has been pointed out by Baratz and Jaffe [BARA-83], there may exist some networks in which that condition cannot be achieved. The following two theorems extend Hagouel's result, as well as Baratz and Jaffe's, by postulating that the shortest path between two nodes, assuming a flat organization, may or may not be fully contained within a 1-cluster common to both nodes.

In the following, we shall refer to the paths obtained on the basis of routing tables structured according to the Kamoun-Kleinrock scheme as *K-paths*. A K-path between two nodes that belong to the same j -cluster must be contained within that cluster in its entirety.

Theorem 1: Consider a two-level K-network in which all 1-clusters have diameters* less than or equal to d node-to-node hops. If w is the length in node-to-node hops of the shortest K-path traversed between two different nodes in the K-network, and w_{opt} is the length of the flat path between the same nodes given a flat network organization, then

$$r = \frac{w}{w_{opt}} \leq 1 + \frac{d}{2} \quad (\text{where } d \geq 1) \quad (1)$$

Proof: If the two nodes (a and b) belong to the same 1-cluster, the K-path between them can be as long as d , the largest diameter of a 1-cluster. On the other hand, as shown in Figure 3, the two nodes could be connected to the same boundary node of an adjacent 1-cluster; hence we obtain the following:

$$r \leq \frac{d}{2} \leq 1 + \frac{d}{2} \quad (2)$$

Now let us assume that the two nodes a and b are located in different 1-clusters ($C(a)$ and $C(b)$, respectively), and let us consider Figure 4. The minimum K-path obtained between a and b equals $x_1 \cup x_2 \cup x_3$, where (a) x_1 is the (minimum) path traversed within $C(a)$ from Node a to a boundary node; (b) x_2 is the minimum K-path between $C(a)$ and $C(b)$; and (c) x_3 is the minimum path between the boundary node reached at $C(b)$ and Node b . Similarly, the shortest

* The diameter of a network is the length of the longest minimum route between any two of its nodes.

path between the same pair of nodes, given a flat network organization, equals $y_1 \cup y_2 \cup y_3$ (as shown in Figure 4). Hence:

$$r = \frac{|x_1| + |x_2| + |x_3|}{|y_1| + |y_2| + |y_3|} \quad (3)$$

According to the Kamoun-Kleinrock scheme, in a 2-level K-network each node knows the shortest distance to every 1-cluster. Hence, the K-path from any Node a in $C(a)$ to a boundary node in $C(b)$ (e.g., BN_b) has shortest length, which means that

$$|y_1| + |y_2| + |y_3| \geq |x_1| + |x_2| \quad (4)$$

simply because BN_b must be closer (in hops) to a than b (Figure 4). Using the inequality of (4) in (3), we obtain:

$$r \leq 1 + \frac{|x_3|}{|y_1| + |y_2| + |y_3|} \quad (5)$$

In the worst case, $|x_3|$, the length of the minimum path between the boundary node reached at $C(b)$ and Node b , can be as long as d . On the other hand, $|y_3|$ may be as small as 0 or 1, depending on whether the destination node, b , is a boundary node or not.

Let Node b be a boundary node, i.e., $|y_3| = 0$. Then equation (5) becomes

$$r \leq 1 + \frac{d}{|y_1| + |y_2|} \quad (6)$$

Because nodes a and b lie in different 1-clusters and no cluster overlap may occur, $|y_2| \geq 1$; the worst case occurs when $|y_2| = 1$. If Node a is a boundary node ($|y_1| = 0$) then Node a must be adjacent to Node b ($|x_1| = 0$); because every node must know its neighbors, this implies that $r = 1 \leq 1 + d/2$. If Node a is not a boundary node, then $|y_1| > 0$, and the worst case is obtained when $|y_1| + |y_2| = 2$ (i.e., either when a is a boundary node and there is an intermediate node between a and b , or when both $|y_1|$ and $|y_2|$ equal 1). For this case Equation (1) holds.

Now assume that Node b is not a boundary node, i.e., $|y_3| \geq 1$. Then Equation (5) becomes

$$r \leq 1 + \frac{d}{|y_1| + |y_2| + 1} \quad (7)$$

Again, because a and b lie in different clusters, $|y_1| + |y_2| \geq 1$, and the worst case satisfies Equation (1). \square




FIGURE 3




FIGURE 4

Theorem 2: Consider an m -level K-network in which all 1-clusters have diameters of lengths less than or equal to d node-to-node hops. Under stationary conditions, if w is the length (in node-to-node hops) of the K-path traversed between two nodes in different k -clusters of the K-network ($k \leq m$), and w_{opt} is the length (in node-to-node hops) of the flat path between the same nodes, then

$$r_k = \frac{w}{w_{opt}} \leq 2^{k-2} \left(2 + \frac{d}{2} \right) - 1 \quad \text{for } 2 \leq k \leq m. \quad (8)$$

Proof: By induction on k .

Let r_i ($i = 1, 2, \dots, k$) denote the ratio of w over w_{opt} for an i -level K-network. For $i = 1$, r_1 must be one since 0-level clusters correspond to the nodes themselves. From theorem 1, we know that $r_2 \leq 1 + \frac{d}{2} = 2^0 \left(2 + \frac{d}{2} \right) - 1$.

Now assume that, at level $k = n$, r_n obeys the inequality $r_n \leq 2^{n-2} \left(2 + \frac{d}{2} \right) - 1$, and postulate a K-network of $n + 1$ levels. Figure 5 shows the minimum K-path between an arbitrary pair of nodes a and b that belong to different n -clusters. It also shows a minimum path when no clustering is assumed. The minimum K-path between a and b equals $x_1 \cup x_2 \cup x_3$; similarly, the shortest path between the same pair of nodes, given a flat network organization equals $y_1 \cup y_2 \cup y_3$. Hence:

$$r_{n+1} = \frac{|x_1| + |x_2| + |x_3|}{|y_1| + |y_2| + |y_3|} \quad (9)$$

Let α be an shortest path between Boundary Node BN_b and b when no clustering is assumed; because α is the shortest, it follows that

$$|\alpha| \leq |x_1| + |x_2| + w_{opt} \quad (10)$$

From our inductive assumption, we know that

$$|x_3| \leq |\alpha| \left(2^{n-2} \left(2 + \frac{d}{2} \right) - 1 \right) \quad (11)$$

Using the same arguments as in Theorem 1, we obtain $|x_1| + |x_2| \leq w_{opt}$. Substituting this inequality, and (10) and (11) in (9), we obtain

$$r_{n+1} \leq 2^{n-1} \left(2 + \frac{d}{2} \right) - 1 \quad (12)$$

and the theorem follows by induction. \square

3.1.2 G-Networks

Consider a connected G-network. Refer to the paths obtained according to the cluster-level route tables of GRNs as *G-paths*. A G-path between two nodes in the same k -cluster must be contained fully within that cluster. The following results specify the worst-case ratio between optimum paths obtained in flat networks and optimum G-paths.

Theorem 3: Consider a G-network of two levels (i.e., nodes and 1-clusters) in which all 1-clusters have diameters less than or equal to d node-to-node hops. If w is the length in node-to-node hops of the shortest G-path traversed between two different nodes in the G-network, and w_{opt} is the length of the flat path between the same nodes, then

$$r = \frac{w}{w_{opt}} \leq d + 1 \quad (13)$$

Proof: As was the case in the proof of Theorem 1, if the two nodes (a and b) belong to the same 1-cluster, then

$$r \leq \frac{d}{2} \leq d + 1 \quad (14)$$

Now assume that the two nodes a and b are located in 1-clusters ($C(a)$ and $C(b)$, respectively) and consider Figure 4 again. The minimum G-path obtained between a and b equals $x_1 \cup x_2 \cup x_3$; the shortest path between the same pair of nodes, given a flat network organization, equals $y_1 \cup y_2 \cup y_3$. Hence, r equals the RHS of Equation (3).

Intercluster routing is based on the GRTs that only GRNs update and distribute among the nodes in their own clusters. Hence, a given G-path between nodes in different 1-clusters is selected based only on two things: (1) the length in GRN-to-GRN hops of the intercluster path from the originating node's GRN to the destination node's GRN; (2) the length of the path from the originating node to the nearest boundary node (in terms of node-to-node hops) in the node's 1-cluster that leads to the next GRN in the path to the destination GRN. Because w is a shortest G-path, c_{x_1} , the number of 1-clusters traversed in x_2 , must be less than or equal to c_{y_1} , the number of clusters traversed in y_2 . Let $c_{y_1} = K - 1$, then $c_{x_1} \leq K - 1$.

Since the number of hops in each cluster included in y_2 must be at least one, it follows that $|y_2| \geq K$ (Figure 4). On the other hand, from Figure 4 we observe that

$$|x_2| = K + \sum_{i=1}^{K-1} l_i \quad (15)$$

Since by assumption all the clusters in the G-network have diameters less than or equal to d , we obtain that $|x_2| \leq K + (K-1)d$. Because Node a does not know the length of the path along which the message will travel within $C(b)$ and a 1-cluster can have a diameter as long as d , $|x_3|$ can be as long as d . Hence:

$$r \leq \frac{|x_1| + K(1+d)}{|y_1| + |y_3| + K} \quad (16)$$

As it was done for Theorem 1, let Node b be a boundary node, then $|y_3| = 0$. If Node a is a boundary node ($|y_1| = 0$) then $|x_1| = 0$ and the RHS of Equation (16) becomes $1 + d$. If Node a is not a boundary node ($|y_1| > 0$) then the worst case occurs when $|x_1| = d$ and $|y_1| = 1$. With these values Equation (16) becomes

$$r \leq d + \frac{K}{K+1} \quad (17)$$

The RHS of (17) has $d + 1$ as its upper bound when K tends to infinity.

Now assume that Node b is not a boundary node, then $|y_3| \geq 1$. If Node a is a boundary node, then $|y_1| = |x_1| = 0$, and with $|y_3| = 1$ (worst case), Equation (16) becomes

$$r \leq \frac{K(1+d)}{1+K} \quad (18)$$

The RHS of (18) also has $1 + d$ as its upper bound when K tends to infinity. If Node a is not a boundary node, then $|y_1| \geq 1$ and $|x_1| \leq d$. Hence:

$$r \leq \frac{d(1+K) + K}{2+K} \quad (19)$$

Again, the RHS of (19) has $1 + d$ as its upper bound as K tends to infinity, and the theorem follows. \square

Corollary 1: Assume an m -level G-network in which all 1-clusters have diameters of lengths less than or equal to d node-to-node hops. Under stationary conditions, if w is the weight of the G-path traversed between two nodes in two different k -clusters of the G-network ($k \leq m$), and w_{opt} is the weight of the flat path between the same nodes, then

$$r_k = \frac{w}{w_{opt}} \leq 2^{k-2}(2+d) - 1 \quad \text{for } 2 \leq k \leq m. \quad (20)$$

Proof: This corollary follows immediately from Theorem 2 and Theorem 3. The proof is by induction on k as in Theorem 2, but with $r_2 \leq d + 1$ (which follows from Theorem 3). \square

3.2 OVERLAPPING

Overlapping of clusters may offer a number of advantages in terms of the PRNET's vulnerability in case of resource failure. Furthermore, it could permit the smooth transition of a node from one cluster to another, as well as the routing of messages to nodes even after the failure of some GRNs or the partition of 1-clusters. Here, we analyze whether overlapping also makes G-network paths and K-network paths significantly shorter.

Corollary 2: Consider an m -level K-network in which all 1-clusters have diameters of lengths less than or equal to d node-to-node hops. Assume further that all adjacent 1-clusters overlap in at least one boundary node. Under stationary conditions, the ratio of the length of the K-path between two nodes in two different k -clusters of the network and the flat path between the same nodes, if we assume a flat network organization, is bounded by

$$r_k = \frac{w}{w_{opt}} \leq 2^{k-2} \left(2 + \frac{d}{2} \right) - 1 \quad \text{for } 2 \leq k \leq m. \quad (21)$$

Proof: The proof of this corollary follows exactly the form of theorems 1 and 2. For $k = 1$, the proof is exactly the same as in Theorem 1, except that in this case boundary nodes must belong to all the 1-clusters that they join. \square

Hence, overlapping of 1-clusters does not reduce the worst-case ratio of K-path lengths to flat path lengths.

Corollary 3: Consider a two-level G-network in which all 1-clusters have diameters less than or equal to d node-to-node hops. Assume further that overlapping must occur among adjacent 1-clusters. Under stationary conditions, the ratio of the length of the G-path between two nodes of the G-network and the flat path between the same nodes is bounded by

$$r = \frac{w}{w_{opt}} \leq d \quad (22)$$

Proof: The proof of this corollary follows exactly the form of Theorem 3. In this case, however, if the two nodes belong to the same 1-cluster, the G-path between them can be as long as d , while the two nodes could be connected to another node c in another cluster through a path of four hops, as shown in Figure 6. Hence, $r \leq d/4 \leq d$. Another difference in the proof is that the lengths of the links between adjacent 1-clusters equal 0 because adjacent 1-clusters must overlap. Therefore;

$$|x_2| = \sum_{i=1}^{K-1} l_i \leq (K-1)d \quad (23)$$

Using the above expression for $|x_2|$, following the same procedure as in Theorem 3 and considering the limit as $K \rightarrow \infty$, we obtain the result in (22). \square

Corollary 4: Consider an m -level G-network in which all 1-clusters have diameters of lengths less than or equal to d node-to-node hops and in which adjacent 1-clusters overlap in their boundary nodes. Under stationary conditions, the ratio of the length of the G-path between two nodes in two different k -clusters of the network ($k \leq m$) and the flat path between the same nodes, if we assume no clustering, is bounded by

$$r_k = \frac{w}{w_{opt}} \leq 2^{k-2}(1+d) - 1 \quad \text{for } 2 \leq k \leq m. \quad (24)$$

Proof: The proof of this corollary follows directly from Theorem 2 and Corollary 3. The proof is by induction on k as in Theorem 2, but with $r_2 \leq d$ (which follows from Corollary 3). \square

Equations (20) and (24) show that there is some improvement in the worst-case ratio of G-paths to flat paths when 1-clusters overlap.

3.3 TIME OF CONVERGENCE

To provide an unbiased comparison between hierarchical and flat routing schemes, instead of considering any particular update algorithm, we will assume that the *best-possible update algorithm* is used in both cases. The best possible update algorithm (which we shall call the *BPU algorithm*) for a given network (flat or hierarchical) converges as fast as possible and with a minimum number of update messages. For a network of V nodes and diameter D , such an algorithm would require D synchronous update cycles to converge to a stable state [SCHW-80] in the worst case (e.g., when *all* routing tables must be updated). The fact that a BPU algorithm takes D cycles to complete implies that an update message must be forwarded through a chain of D internodal hops from the beginning to the end of the update procedure.




FIGURE 5




FIGURE 6

A BPU algorithm can be either periodical or event-driven. A periodical BPU algorithm is one in which update messages are sent out by each node at regular intervals; in an event-driven BPU algorithm updates are sent out as a result of a topological change by those nodes affected by such a change. By assuming reliable transmissions and a constant propagation delay over the radio links, synchronous operation of a BPU algorithm is obtained. The number of synchronous update cycles times the longest propagation delay for each update provides an upper bound on the convergence time of the algorithm [JOHN-83].

In the case of K-networks and flat PRNETs, an event driven BPU algorithm would create too many collisions in the radio channel; hence, we will assume a periodic BPU algorithm in such cases. For the case of G-networks, we will assume that a periodical BPU algorithm is used to update *NRTs*, and that an event-driven BPU algorithm is used to update *GRTs*, and assume synchronous operation of such algorithms to analyze the worst case. This strategy is feasible because of the relatively small size of 1-clusters and the relatively few *GRNs* of a G-network.

The results of this section provide a lower bound for the worst-case convergence times of routing algorithms in flat and hierarchical PRNETs (G-networks and K-networks). We use such bounds *solely* to assess the relative benefits of hierarchical and flat *network organizations*; worst-case convergence times for specific algorithms are presented elsewhere [GARC-84].

Throughout this subsection we assume that no transmission errors occur in the radio channel and that all radio links are bidirectional.

3.3.1 Updates Affecting *NRTs*

As we have stated, the update of *NRTs* is carried out within a 1-cluster just as in a PRNET with a flat organization. The following proposition specifies the worst-case convergence time after failure in flat PRNETs.

Proposition 1: Consider a flat PRNET with diameter D . Assume that a topological change occurs and that there are no more after that. Then, the time required by the PRNET to converge to a stable state after the topological change has been detected is bounded by $[D \cdot T_p]$ under the BPU algorithm. T_p is the time between the transmission of two consecutive periodic updates by the same node.

Proof: The nodes of a flat PRNET transmit their updates periodically every T_p seconds; nodes in radio connectivity must access the channel at different times to avoid collisions. In the worst case, after a Node a has sent its update corresponding to Cycle j , it will not be able to receive and process the update for Cycle j from a neighbor b before it has to send its update for Cycle $j + 1$. Hence, each update cycle will take T_p seconds at each node and the proposition

holds because the BPU algorithm takes D synchronous update cycles to converge. \square

It follows from the above that, in a PRNET organized as a G-network but using a periodic BPU algorithm to update NRT s, the maximum convergence time after a topological change that does not affect the connectivity among GRN s is $[d \cdot T_p]$, where d is the maximum diameter of a 1-cluster of the G-network and T_p is the time between two NLU transmissions from the same node.

3.3.2 Updates Affecting GRTs

If the proposed hierarchical routing scheme is implemented, the contents of GRT s need to be changed only after the following two cases of topological change: (1) changes on links or boundary nodes joining two or more clusters; (2) the addition, deletion, or partition of a cluster. We will consider just the first case, as it does not involve cluster reconstitution mechanisms that are highly dependent on the type of overlapping among clusters, the treatment of which lies beyond the scope of this paper.

The following proposition and theorem show the worst-case convergence time after a cluster-level topological change for K-networks and G-networks.

Proposition 2: Consider a PRNET that uses a periodic BPU algorithm and whose nodes are organized as an m -level K-network. Assume that a topological change occurs, and that there are no more after that. Then, under synchronous operation, the time t_c required by the network to converge to a stable state after such a topological change is detected is bounded by $t_c \leq D T_p$, where D is the diameter of the same network when a flat organization is assumed, and T_p is the time between the transmission of two consecutive periodic updates by the same node.

Proof: According to the scheme proposed by Kamoun and Kleinrock [KAMO-76], the network is organized into nonoverlapping clusters, and the nodes participate as peers in the updating of their routing tables. An example of the structure of such routing tables is depicted in Figure 2 (b). Each routing table entry contains the length in node-to-node hops of the minimum path to either a given node in the same 1-cluster or the boundary node of a remote cluster [KAMO-76]. Hence, the failure or addition of a single node-to-node link or a single node could indeed affect the distance from a node to a distant cluster (e.g., in Figure 2 (a), the failure of link (d,e) would cause the distance from Node 6 to Cluster C.3 to increase by one).

In the worst case, the distance in node-to-node hops from the node that detected a topological change affecting intercluster connectivity and a given node could be as long as D , the diameter of the network with flat organization. Since updates among nodes are transmitted periodically every T_p seconds by any given node, this proposition follows from Proposition 1. \square

Theorem 4: Consider an m -level G-network whose 1-clusters are *disjoint* and with diameters shorter than or equal to d node-to-node hops. Assume that a periodic BPU algorithm is used to update *NRTs* and an event-driven BPU algorithm is used to update *GRTs*. Furthermore, assume that a change in intercluster links or boundary nodes occurs that affects the connectivity among *GRNs* in the same $(k + 1)$ -cluster, and that there are no more topological changes afterwards. Let T_E be the longest propagation time for a *CLU* forwarded between two adjacent nodes, and T_p be the time between two consecutive *NLU* transmissions. Then, the time required for that PRNET to converge to a stable state after such topological change is bounded by

$$t_c \leq [(1 + d)^m - (1 - d)] T_E + d T_p \quad (25)$$

Proof: For a *GRN* to note that its logical connectivity with another adjacent *GRN* has changed, it must be notified by any of the boundary nodes of its 1-cluster. This could take as long as $d T_E$ seconds because a boundary node can be as many as d hops away from its *GRN*. Only after this time has elapsed can the cluster-level update procedure start among *GRNs*.

Let D_k be the *maximum* diameter of a k -cluster measured in node-to-node hops. We have assumed that the maximum diameter of a k -cluster is d $(k - 1)$ -cluster-to- $(k - 1)$ -cluster hops; hence, a path of length D in the k -cluster will include $[d + 1]$ $(k - 1)$ -clusters connected by d links between boundary nodes. Therefore, D_k can be expressed recursively as follows:

$$D_k = (1 + d) D_{k-1} + d \quad (26)$$

$$D_1 = d$$

By induction on k , the solution to the above recurrence relation can be shown to be $D_k = (1 + d)^k - 1$. Hence, a packet sent across an m -cluster may have to traverse as many as $(1 + d)^m - 1$ node-to-node hops within the cluster. If a BPU algorithm is used and *all* *GRNs* had to update their *GRTs* (worst case), the entire diameter of the only m -cluster in the PRNET may have to be traversed (from the *GRN* that starts the cluster-level update procedure to the last *GRN* that receives a *CLU* during the last cycle of the procedure). Hence, it follows that the completion of the cluster-level update procedure among *GRNs* may account for as many as $[(1 + d)^m - 1] T_E$ seconds in the worst case.

Finally, each *GRN* has to distribute the resulting updates in its *GRT* to the rest of the nodes in its 1-cluster. Since such updates are distributed by means of *NLU*s, it follows from Proposition 1 that this process can take as long as $[d \cdot T_p]$ seconds. Hence, as many as $[d \cdot T_p]$ seconds may elapse after all *GRN*s have updated their *GRT*s before all nodes in all clusters have consistent *GRT*s.

Adding up the three terms obtained above, we get the result in Equation (25). \square

4. NETWORK STRUCTURE OPTIMIZATION

In the previous section, we looked at two hierarchical network organization strategies (G-networks and K-networks) and obtained upper bounds on the convergence time and shortest path lengths in PRNETs organized according to those two schemes. Such bounds depend on the number of cluster levels and the maximum diameter of the clusters (which in turn depends on the number of elements in a cluster). It is clear that shorter cluster diameters and fewer cluster levels will provide shorter convergence times and routes. However, the number of elements c in a k -cluster (i.e., $(k - 1)$ -clusters) and the number of cluster levels, m , are related to each other for $m \geq 2$. In other words, for a fixed number of nodes, the smaller a cluster is in an m -level PRNET ($m \geq 2$), the more clusters are needed and (potentially) the more cluster levels there must be to obtain the same upper bounds on path lengths in G-networks and K-networks. In this section we establish a relation between the number of cluster levels and the size of clusters that minimizes the upper bound of optimum G-path and K-path lengths, which are indeed achievable. Previous work on the design of hierarchical networks has focused on minimizing the length of routing tables [KAMO-76; SHAC-84b]. However, while reducing the length of routing tables is important in large networks with hierarchical structures, an optimum table size is not as critical a design objective as reducing the length of the paths traversed in such networks.

To simplify the problem of determining an optimum number of clusters and an optimum cluster size that would minimize the upper bound of G-path lengths we treat the number of cluster levels (m) and the number of elements in a cluster (c) as real numbers. Furthermore, we assume that every k -cluster is formed by exactly c $(k - 1)$ -clusters.

Consider an m -level G-network with nonoverlapping 1-clusters. Let V be the number of nodes in the network and c be the size of every cluster at every level. Because the diameter of a cluster must be less than or equal to $c - 1$, it follows from (20) that

$$r_m \leq 2^{m-2} (c + 1) - 1 \quad (27)$$

Because there is no overlap among 1-clusters and since all clusters at all levels are assumed to have the same size c , it follows that c must equal $V^{1/m}$. Substituting this value of c in (27), and defining R_m as the maximum value that r_m can take, we obtain

$$R_m = 2^{m-2} \left(V^{\frac{1}{m}} + 1 \right) - 1 \quad (28)$$

Equation (28) shows that R_m is continuous for all values of m . Taking the derivative of R_m with respect to m and equating to zero we obtain the following equality:

$$\left[V^{-\frac{1}{m}} + 1 \right] m^2 \frac{\ln 2}{\ln V} = 1 \quad (29)$$

Solving (29) for all real values of m is a rather difficult task. Fortunately, we can simplify the problem significantly by considering large values of V . For large values of V , R_m can be approximated by

$$R_m = 2^{m-2} (c) - 1 ; \quad (30)$$

from which we obtain that the real value of m that minimizes R_m is given by:

$$m_{opt} \approx \left(\frac{\ln(V)}{\ln(2)} \right)^{\frac{1}{2}} ; \quad V \gg 1 \quad (31)$$

With c being $V^{1/m}$, it follows that the length of the routing table size, η , of every node would be

$$\eta = m V^{\frac{1}{m}} \quad (32)$$

Consider now an m -level G -network in which all adjacent 1-clusters must overlap in at least one boundary node. Again, because the diameter of a cluster must be less than or equal to $c - 1$, it follows from (24) that

$$R_m = 2^{m-2} (c) - 1 \quad (33)$$

Because adjacent clusters must overlap and all clusters must be of the same size, it follows that $c \geq V^{1/m} + 1$. Taking the smallest possible cluster size we obtain Equation (28). With c being $V^{1/m} + 1$, it follows that the length of the routing table size, η , of every node would be

$$\eta_{overlap} = m \left[V^{\frac{1}{m}} + 1 \right] \quad (34)$$

As expected, overlapping of 1-clusters requires longer routing tables; however, the overhead imposed by overlapping 1-clusters is in $O(\sqrt{m})$ of m , which is quite small for the values of m expected in real PRNETs.

For the case of K-networks with nonoverlapping 1-clusters, we have that

$$R_m = 2^{m-3} \left(3 + \frac{1}{V^m} \right) - 1 \quad (35)$$

If we make the assumption that $V \gg 3$ we obtain the same results presented above for G-networks.

Equation (31) provides only an approximation of the real optimum value of m when V is very large. Table I shows the value of m_{opt} obtained by solving (29) numerically and by using Equation (31). It is clear from these results that the error incurred with our approximation is relatively small (smaller than 0.2) and always positive. Since we are only interested in integer values of m_{opt} (denoted m_{opt}^*), the RHS of (31) provides a very good approximation of the true optimum value of m .

Table I
OPTIMAL VALUES OF m

V	m_{opt} (numeric solution)	m_{opt} (Eq. (31))
10	1.6339	1.8226
10^2	2.4062	2.5776
10^3	3.009	3.1569
10^4	3.519	3.6452
10^5	3.9679	4.0755
10^6	4.3726	4.4645
10^7	4.7435	4.8222

5. SUMMARY AND DESIGN CONSIDERATIONS

In this paper we have analyzed the performance of a hierarchical routing scheme for large-scale PRNETs relative to that of flat PRNETs and those organized by means of the Kamoun-Kleinrock scheme [KAMO-76]. Our focus was on two principal performance figures: (1) the quality of the routing decisions made by the nodes; (2) the time required for all nodes to update their routing tables after topological changes. Our results on optimum path lengths expand upon previous results by Hagouel [HAGO-83] and Baratz and Jaffe [BARA-83]. While only worst-case conditions were analyzed, we can draw a number of important design considerations as to the structure a hierarchical PRNET should have.

Table II lists the values of m (Equation (31)) and c (i.e., $V^{1/m}$) that would minimize the upper bound of G -path lengths in G -networks with no cluster overlapping for five different values of V . The values of η (Equation (32)) and r_m are also shown. It is interesting to note that, while minimizing the value of η was not our goal, η remains within very reasonable bounds even for very large networks. There is a good reason for preferring a small η , which is the fact that NLU s in G -networks, and all updates in K -networks, contain complete routing tables. In G -networks, however, there exists the alternative of distributing cluster-level updates to simple nodes and boundary nodes using CLU s rather than as part of periodic NLU s (as it has been proposed in this paper). This would mean that GRT s could be of any size, but would increase the traffic in the network. It is clear from the data in Table II that compact clusters, i.e., clusters with short diameters, are desirable in G -networks for achieving short G -paths and short convergence times.

In contrast to our approach, Kamoun and Kleinrock optimized the values for c and m with respect to the value of η for the case of K -networks, which results in $m = \ln(V)$ and $c = e$. Table III shows the values of m , c , η , and r_m (Equation 8) for the same values of V of Table II. Optimizing the organization of a hierarchical network on the basis of the length of its routing tables can potentially result in very long hierarchical paths (as compared with the shortest paths obtained in flat networks) because of the resulting large number of cluster levels.

In the absence of collisions in the channel, the worst-case convergence time of a PRNET with a hierarchical routing scheme based on GRN s (a G -network) is much faster than the worst-case convergence time of a flat PRNET or a K -network. The main reason for this is that, in G -networks, cluster-level updates can be propagated across clusters on an event-driven basis (rather than periodically) by only a few selected nodes. Since the transmission time of a CLU is very short (as it contains only those entries that must be updated in GRN 's GRT s) and the radio channel of a PRNET has a high transfer rate (400 kbps in the DARPA PRNET) the time needed to propagate a CLU between two adjacent nodes (T_E) is very small. In contrast, the time between

periodic updates in flat PRNETs has to be relatively long (e.g., about 7.5 seconds in the DARPA PRNET [WEST-82]).

It is easy to see from propositions 1 and 2 and Theorem 4 that when the best possible update algorithms are assumed for both G-networks and K-networks, the worst-case convergence time of G-networks is indeed much smaller than those of flat PRNETs and K-networks. As an example, assuming that $d = c/2$ (a conservative estimate) and using the entries in Table II, it can be shown that $t_c = O(d T_p)$ in all cases; on the other hand, in real networks it is reasonable to expect that the diameter of the whole PRNET is much longer than the diameter of a 1-cluster ($D \gg d$) specially for large V .

In the worst case, as expected, optimum G-paths and K-paths between two distant nodes can be much longer than the shortest paths obtained with a flat network organization between the same nodes, provided that the routing tables of the flat network were all correct. However, because periodic updating algorithms need to be used in flat PRNETs, it may take many seconds to obtain consistent routing tables after topological changes affecting the connectivity of a large number of nodes. Hence, highly suboptimal routes may be generated in large flat PRNETs.

Since usually $d \geq 2$, the paths that can be obtained with the proposed hierarchical network organization can be longer than those achievable by using the Kamoun-Kleinrock scheme alone, provided that the same number of cluster levels and cluster size have been used in both types of networks. However, the ratio between the two cannot exceed 2. Furthermore, if the optimum structure proposed by Kamoun and Kleinrock (obtained by minimizing η) were used, it could be possible to obtain a lower quality of routing in the K-network than in the corresponding optimum G-network even for the case in which d equals $c - 1$ (see Table II).

As expected, overlapping of 1-clusters in K-networks does not provide any advantages in terms of path lengths. From the results in (22) and (24), it is clear that overlapping of clusters in G-networks provides some advantage from the standpoint of intercluster path length; furthermore, as it was discussed in Section 4, the overhead in routing table size imposed by overlapping 1-clusters is really small (of the order of m). However, the desirability of overlapping should be assessed according to the robustness it would provide to the network and the complexity of the algorithms it would introduce. Shacham [SHAC-84b] discusses the issue of overlapping clusters in more detail.

We can conclude from the foregoing that the proposed G-networks constitute a viable approach to the organization of large PRNETs with mobile nodes; they provide a reasonable compromise between quality of routing decisions and the speed with which routing table updates are propagated. Further research will be necessary to clarify such issues as: (1) their performance (e.g., throughput and end-to-end delay) under average conditions assuming specific routing algorithms; (2) the optimization of such networks with respect to path lengths; (3) the effect of overlapping on the capacity of the PRNET to respond to cluster partitions, creations, and deletions.

Table II
G-NETWORKS

NODES	$m(\text{Eq. 31})$ (integer value)	$\sqrt{1/m}$	$\eta = mc$	$r_m \leq (\text{Eq. 20})$
10^2	3	5	15	$2d + 3$
10^3	3	10	30	$2d + 3$
10^4	4	10	40	$4d + 7$
10^5	4	18	72	$4d + 7$
10^6	5	16	80	$8d + 15$

Table III
K-NETWORKS

NODES	$m = \ln V$	$c = 3$	$\eta = mc$	$r_m \leq (\text{Eq. 8})$
10^2	5	3	15	$4d + 15$
10^3	7	3	21	$16d + 63$
10^4	9	3	27	$64d + 255$
10^5	12	3	36	$512d + 2047$
10^6	14	3	42	$16384d + 32767$

NOTE:

d = largest diameter of any cluster ($\leq c-1$)

η = length of routing table

m = number of cluster levels

c = size of a cluster

6. REFERENCES

- [BAKE-81] D.J. Baker and A. Ephremides, "The Architectural Organization of a Mobile Radio Network via a Distributed Algorithm," *IEEE Transactions on Communications*, Vol. COM-29, No. 11, Nov. 1981, pp. 1694-1701.
- [BAKE-83] D.J. Baker, "Distributed Control of Broadcast Radio Networks with Changing Topologies," *Proceedings IEEE INFOCOM 83*, San Diego, California, April 1983, pp. 49-55.
- [BARA-83] A.E. Baratz and J.M. Jaffe, "Establishing Virtual Circuits in Large Computer Networks," *Proc. IEEE INFOCOM '83*, San Diego, California, April 1983, pp. 311-318.
- [GAFN-81] E.M. Gafni and D.P. Bertsekas, "Distributed Algorithms for Generating Loop-Free Routes with Frequently Changing Topology," *IEEE Transactions on Communications*, Vol. COM-29, No. 1, January 1981, pp. 11-18.
- [GARC-84] J.J. Garcia-Luna-Aceves, "Responsive Routing Algorithms for Packet Radio Networks," in preparation.
- [HAGO-83] J. Hagouel, "Issues in Routing for Large and Dynamic Networks," *IBM Research Report RC 9942 (No. 44055) Communications*, IBM Thomas J. Watson Research Center, Yorktown Heights, New York 10598, April 1983.
- [JOHN-83] M.J. Johnson, "Analysis of Routing Table Update Activity after Resource Failure in a Distributed Computer Network," *Proc. SIGCOMM '83 Symposium*, Austin, Texas, March 1983, pp. 14-20.
- [KAHN-78] R.E. Kahn, S.A. Gronemeyer, J. Burchfield, and R.C. Kunzelman, "Advances in Packet Radio Networks," *Proceedings of the IEEE*, Vol. 66, No. 11, November 1978, 1468-1496.
- [KAMO-76] F. Kamoun, "Design Considerations for Large Computer Communication Networks," *UCLA-ENG-7642*, Computer Science Department, School of Engineering and Applied Science, University of California, Los Angeles, January March 1976.
- [MACG-82] W. MacGregor, J. Westcott, and M. Beeler, "Multiple Control Stations in Packet Radio Networks," *Conference Record of IEEE Military Communications Conference*, Boston, Massachusetts, October 1982, Vol. 3, pp. 10.3.1-10.3.5.
- [NILS-80] A. Nilsson, W. Chou, and C.J. Graff, "A Packet Radio Communication System Architecture in a Mixed Traffic and Dynamic Environment," *Proceedings of Computer Networking Symposium*, pp. 51-66, December 1980.
- [SCHW-80] M. Schwartz, "Routing and Flow Control in Data Networks," *IBM Research Report RC 8353 (No. 36329)*, IBM Thomas J. Watson Research Center, Yorktown Heights, New York, 1980.
- [SHAC-84a] N. Shacham and K. Klemba, "An Architecture for Large Packet Radio Networks and some Implementation Considerations," *SRNTN 11*, Contract No. MDA-903-83-C-0155, SRI Project 5732, SRI International, Menlo Park, California 94025, February 1984.
- [SHAC-84b] N. Shacham, "Organization of Dynamic Radio Network by Overlapping Clusters: Architecture Considerations and Optimization," *Proc. Performance '84*, Paris, France, December 1984.

- [SHAC-84c] N. Shacham, "Hierarchical Routing in Large Dynamic Packet Radio Networks," SRNTN 18, Contract No. MDA-903-83-C-0155, SRI Project 5732, SRI International, Menlo Park, California 94025. Also in Proc. of the 18th Hawaii International Conference on System Sciences, Honolulu, Hawaii, January 1985.
- [WEST-82] J. Westcott and J. Jubin, "A Distributed Routing Design for a Broadcast Environment," Conference Record of IEEE Military Communications Conference, Boston, Massachusetts, October 1982, Vol. 3, pp. 10.4.4-10.4.5.